



CERTIFIED COPY OF PRIORITY DOCUMENT

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name by which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., p.l.c. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.


Signed

Dated 15 May 2001



Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference

N00/0407/GB

2. Patent application number

(The Patent Office will fill in this part)

0013571.5

6 JUN 2000

3. Full name, address and postcode of the or of each applicant (*underline all surnames*)

Power X Limited
Stafford Court
145 Washway Road
Sale
Cheshire M33 7PE

Patents ADP number (*if you know it*)

If the applicant is a corporate body, give the country/state of its incorporation

GREAT BRITAIN

6803233002

4. Title of the invention

SWITCHING SYSTEM

5. Name of your agent (*if you have one*)

"Address for service" in the United Kingdom to which all correspondence should be sent (*including the postcode*)

McNeight & Lawrence
Regent House, Heaton Lane
Stockport, Cheshire SK4 1BS

Patents ADP number (*if you know it*)

0001115001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (*if you know it*) the or each application number

Country

Priority application number
(*if you know it*)

Date of filing
(*day / month / year*)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(*day / month / year*)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (*Answer 'Yes' if:*

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description	18
Claim(s)	5
Abstract	871
Drawing(s)	7 + 7

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and search (*Patents Form 9/77*) 1

Request for substantive examination (*Patents Form 10/77*)

Any other documents
(please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature

Date

McNeight & Lawrence

5 June 2000

12. Name and daytime telephone number of person to contact in the United Kingdom
- J G Lawrence
0161 480 6394

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.*
- Write your answers in capital letters using black ink or you may type them.*
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.*
- Once you have filled in the form you must remember to sign and date it.*
- For details of the fee and ways to pay please contact the Patent Office.*

Switching System

Field of the invention

The present invention relates to devices and apparatus for data switching. One example of the use of the present invention is to provide high bandwidth
5 interconnection within systems in which two or more processors share memory.

Background of the Invention

The ever-expanding requirements for processing-intensive computer applications are driving the market to produce systems of ever-greater power.
10 Unfortunately, improvements in processor technology, though impressive, are insufficient to satisfy all of this demand.

One alternative possibility for creating a system with increased power is to operate several closely coupled processing nodes in tandem. Though each node operates in its own local memory space, the close coupling necessitates
15 a degree of memory sharing. This shared memory can be implemented as a single central copy, or (more typically) replicated and distributed in the nodes' local memory. Either way this gives rise to the need for a high bandwidth inter-node communication system, in the former case to provide access to the central memory, and in the latter case to ensure that the distributed copies are
20 kept coherent.

A node generating traffic through this communication system will frequently require a reply to its request before processing can continue. Thus, either the node must suspend processing, or (where possible) it must switch to another task which is not so stalled – either option will cost overall performance. Low
25 latency in the inter-node communication system is therefore a prime requirement to minimize such loss.

In data communications systems, cell loss can be handled by higher layers in the protocol stack and can therefore be tolerated. By contrast, cell loss in processor interconnect systems is generally unacceptable due to the stalled requesting process, yet such systems typically operate with a minimum of
 5 protocol layers in order to keep down system latency. The physical layer must therefore implement a reliable delivery protocol in hardware.

In GB 9828144.7 (filed on 22 December 1998), we proposed a data switching apparatus which possesses inherent attributes of high bandwidth, scalability, low latency, small physical volume and low cost. Only limited details of this
 10 technology have so far been made publicly available. It is illustrated in Fig. 1.

A switching system employs a number $n+1$ of routers, which may be directional. The information transmission aspect of the respective routers is expressed as "ingress routers" $ITM_0, ITM_1, \dots, ITM_n$. The information receiving aspect of the routers is expressed as the $n+1$ "egress routers" $ETM_0, ETM_1, \dots, ETM_n$. Each router receives information from one or more data sources (e.g. a set of processors constituting a "node"), e.g. ingress router ITM_0 receives information from $m+1$ data sources $ILE_{00}, \dots, ILE_{0m}$. Similarly, each egress router sends information to one or more data outputs, e.g. egress router ETM_0 sends information to data sources $ELE_{00}, \dots, ELE_{0m}$. The master device SC and
 20 matrix device(s) SW constitute the central interconnect fabric (CIF). Cells for transmission through the matrix SW are of equal length, and are each associated with a priority level. Each ingress router maintains, for each egress router and for each priority level, a respective "virtual output queue" of cells of that priority level for transmission to that egress router when the matrix device
 25 SW connects that ingress router to that egress router. Each ingress router sends connection requests to the master device SC. The master device SC determines which ingress and egress routers to connect by a first arbitration process. Each ingress router, having been informed of which egress router it will be connected to, performs a second arbitration to determine which priority

level of cell it will transmit to that egress router, and having determined the priority level, transmits the head of the virtual output queue for that priority level and that egress router to the matrix SW via the serial links to arrive at the same time as connection information sent directly from the master. In
5 practice, the latter is significantly quicker than the former, and has to be artificially delayed in order to match the latency of the path via the router. In summary, the above system uses a memoryless fabric with all congestion buffering in the routers.

10 Summary of the present invention

The present invention aims to provide a new and useful data switching device and method.

In general terms, the present invention proposes that the switching matrix itself maintains a (e.g. short) queue of cells which are to be transmitted.
15 Each of these queues corresponds to one of the virtual output queues stored by the ingress routers, and indeed the cells stored in the switching matrix are replicated from the first cells queuing in the respective virtual output queues. Thus, when it is determined that a connection is to be made between a given input and output of the switching matrix, a cell suitable for transmission along
20 that connection is already available to the switching matrix. It is not necessary to obtain it from an ingress router.

Specifically, in a first aspect the present invention provides a data switching device having a plurality of ingress routers, a plurality of egress routers, a switching matrix and a connection controller,
25 the switching matrix having input ports connected to respective said ingress routers and output ports connected to respective said egress routers, and controlled by the controller to form connections between pairs of input and output ports;

each ingress router including one or more virtual output queues for each egress router, each virtual output queue being arranged to store fixed length cells having a header defining the egress router to be used in the switching matrix connection;

- 5 each input port of the switching matrix including for each virtual output queue in the ingress router connected to that input port a respective head of queue buffer, each buffer being arranged to receive a replication of at least one cell in the corresponding virtual output queue;

- 10 wherein, upon the switching matrix forming a connection between a given input port and output port, the switching matrix transmits from that input port to that output port a cell from one of the one or more corresponding head of queue buffers, and

- 15 upon error free receipt by an egress router of a cell from one of the virtual output queues of one of the ingress routers, a receipt signal is transmitted to that ingress router, that ingress router storing the cell until receiving the receipt signal.

- 20 Upon receipt of a new cell by one of the ingress routers, the cell is stored in one of the virtual output queue(s) of the ingress router corresponding to the egress router for the cell. Each of the virtual output queues and the head of queue buffers may be a first-in-first-out (FIFO) queue, and the head of queue buffer may replicate the first few entries of the virtual output buffer. This may be achieved, for example, by the ingress router, when it receives a new cell and writes it into a virtual output queue, also writing it to the corresponding head of queue buffer, if that buffer has space. If not, the cell
- 25 may be stored, and written to the head of queue buffer when that buffer has space for it.

Thus, the virtual output queues is segregated into two areas, a first area containing cells waiting for replication to the corresponding head of

queue buffer, and a second area containing cells replicated to the head of queue buffer.

For example, one way of determining whether the head of queue buffer has space is to maintain a credit count, indicative of the number of free cells of the corresponding head of queue buffer. When a new cell is being written into a virtual output queue, and the credit count of the corresponding head of queue buffer is not zero, a replication of the cell can be transmitted to that head of queue buffer, and the credit count is decreased by one. Upon the controller causing a connection to switch a cell of an ingress router through the switching matrix, a connection grant signal is transmitted to that ingress router, and increments the credit count by one. Upon determining that there is at least one cell in the first area of a given virtual output queue, and that the number of free cells of the corresponding head of queue buffer is not zero, a replication of at least one cell in the first area is transmitted to that head of queue buffer.

At an appropriate time, e.g. when the ingress router is satisfied that the head of queue buffer replicates the front of the corresponding virtual output queue, the ingress router may transmit a connection request to the controller.

Having received more than one connection request, the controller decides which to satisfy. To begin with, the controller may determine whether any given one of the received requests (e.g. among those requests in relation to cells at the front of one of the the head of queue buffers) can be satisfied without making it impossible to satisfy any of the other received requests. In this case, the controller causes that given request to be satisfied: the cell which is the subject of the request is transmitted. Otherwise (i.e. if at least two requests conflict), the controller may perform an arbitration to decide which to satisfy, e.g. according to known techniques.

As mentioned above, there may be more than one virtual output queue for each pair of ingress and egress routers. For example, the cells may be of different "types", e.g. priority levels, with a different virtual output queue (and thus a different head of line queue) for each type. In this case, the controller
5 may determine, in tandem with which pairs of input and output ports will be connected, the priority level of the cell to be transmitted between them, and transmit that information to the switching matrix, so that a cell is transmitted from the queue having that priority level and the corresponding pair of input and output ports. The determination of which priority level to transmit may be
10 performed by arbitration (e.g. predetermined rules) according to known techniques.

In a second aspect, the invention provides a method of operating a data switching device having a plurality of ingress routers, a plurality of egress routers, a switching matrix and a connection controller, and
15 the switching matrix having input ports connected to respective said ingress routers and output ports connected to respective said egress routers, and controlled by the controller to form connections between pairs of input and output ports;

the method comprising the steps of:

20 maintaining at each ingress router one or more virtual output queues for each egress router, each virtual output queue being arranged to store fixed length cells having a header defining the egress router to be used in the switching matrix connection;

maintaining at each input port of the switching matrix for each virtual
25 output queue in the ingress router connected to that input port a respective head of queue buffer, each buffer being arranged to receive a replication of at least one cell in the corresponding virtual output queue;

upon the switching matrix forming a connection between a given input port and output port, transmitting from that input port to that output port a cell from one of the one or more corresponding head of queue buffers; and

5 upon error free receipt by an egress router of a cell from one of the virtual output queues of one of the ingress routers, transmitting a receipt signal to that ingress router, that ingress router storing the cell until receiving the receipt signal.

Brief description of the drawings

10 An embodiment of the invention will now be described for the sake of example only with reference to the Figures in which:

Fig. 1 shows the system of GB 9828144.7, and of an embodiment of the present invention;

Fig. 2 shows schematically an embodiment of the present invention;

15 Fig. 3 shows schematically processing in the embodiment of Fig. 2;

Fig. 4 illustrates the delays caused by the serial links in the embodiment of Fig. 2;

Fig. 5 illustrates processing according to the present invention in the case that an arbitration is not required;

20 Fig. 6 illustrates processing according to the present invention in the case the an arbitration is required;

Fig. 7 illustrates the average latency according to the present invention, as compared to other switching techniques; and

25 Fig. 8 illustrates processing according to the present invention including a confirmation that a cell has been correctly transmitted.

Detailed description of Embodiments

The embodiment of the present invention described herein is a development of the system described above with reference to Figs. 1, with further reductions to the latency and improvements in fault tolerance. The embodiment is illustrated in Fig. 2, which shows a system having a number (up to 16) of multi-processor nodes 1, 3,...31. Each node contains a router device 33, 35, ...53. The router device provides the interface (both receiving and transmitting information) between each processing node and a central interconnect fabric 57.

The fabric 57 is organised as two independent channels with separate power and clock domains. Each channel consists of a single master and several matrix devices, with the number of matrix devices determining the aggregate bandwidth of the fabric. The router of each node of the multiprocessor system connects into the fabric through an array of high-speed serial links operating over cables. As in the known system described above in relation to Fig. 1, the present embodiment contains 3 types of device: router devices which provide the interface between the interconnect and a processing nodes, a master device (controller) which provides the scheduling and arbitration function in the fabric and one or more matrix devices which provide the crossbar function. The transmission and reception of the nodes along a single one of the channels conforms to the structure explained above and shown in Fig. 1. That is, the routers may be bi-directional routers, in which the data input and output functions may be regarded as ingress and egress routers, and communicate (over one channel) using a master (controller) and cyclic switching matrix.

Under normal failure-free conditions, messages are routed through either of the two channels to balance traffic through the embodiment. When one channel has failed, the other channel is capable of carrying all traffic at a

reduced overall bandwidth. Reliable port-to-port message delivery is ensured through the support of a node-to-node ack/nack protocol whereby every message delivery through the interconnect is reported back to the source node and any corrupted messages are automatically retransmitted.

5 The present embodiment incorporates the following changes from the system described above:

- The master is given control of the selection of the class of each message in order to take over the router arbitration function.

- The matrix maintains a limited store of messages to allow
10 immediate reaction to connections generated by the master, without reference to the router. This is achieved by keeping a set of head of queue (HOQ) buffers in the matrix, one HOQ for each combination of source port, destination port, and message class. There is a 1:1 correspondence between router buffers (VOQs) and matrix HOQs to avoid this additional storage
15 introducing any head of line blocking.

- Under low load situations when message buffers are empty, they can be bypassed to achieve the minimum possible latency.

As discussed below, with these enhancements, the matrix is able to immediately action connections received from the master, without waiting for
20 any actions by the router: all post-master-arbitration router actions are effectively removed from the critical path, resulting in an overall port-to-port latency as low as 55ns.

Fig. 3 shows a logical view of the operation of one ingress and egress port of one channel. When a new message arrives in the ingress router 60, the class and destination are extracted from the header and the message is appended
25 to the appropriate VOQ 62. The message is held in the VOQ 62 until its error-

free receipt has later been acknowledged by the egress router 64. A plurality of matrix devices support a matrix 66, controlled by a controller (master) 68. If the corresponding HOQ in the matrix devices is not full, then at the same time as writing to the VOQ a copy of the message is forwarded to the matrix devices and the router sends a connection_request to the master 68 informing it of the destination and class of the new message. The master maintains a buffer of requests it has not yet satisfied.

On receipt of the connection_request the master 68 immediately signals to the matrix devices which HOQ buffer is to receive the arriving message. This feature is necessary (in this embodiment) since each matrix device receives a different portion of the message, so typically only one of the devices would be capable of decoding this information from the message header.

The master 68 then arbitrates between competing requests and later issues a set of connections to be established by the matrix 66, one of which is for the message under consideration. The master 68 also informs the matrix of the class of message to be transferred for each connection, to identify the specific HOQ containing each message. The matrix 66 can therefore create the connection and forward the message as soon as this information arrives from the master 68. The egress router 64 is sent data_valid to indicate the arrival of the message from the matrix 66 and the ingress router 60 is sent connection_grant to indicate that a message has been forwarded from the Matrix.

When the egress router 64 receives a message it checks the message CRC field and forwards a response (ack indicating correct receipt, otherwise nack) to the originating ingress router. The egress router 64 abandons failing messages, and queues good messages in the appropriate egress queue EQ for that class from where they are transmitted to the node.

The primary means by which latency is reduced compared to the known system discussed above in relation to Fig. 1, is by inclusion of the HOQ buffers, which as explained above remove the path via the connection_grant, ingress router arbitration, and serial links to the matrix from the critical path
 5 on the arbitration of a connection by the master.

Under low load situations where a message arrives at an empty VOQ, it will be passed on to the HOQ and the connection_request will be generated simultaneously with writing it into the VOQ. This avoids the overhead of a buffer write and read.

10 When the master 68 receives the connection_request, if there are no competing requests for either the ingress or egress, the master 68 can bypass the normal arbitration algorithm and generate an immediate connection to the matrix 66. This replaces the normal specification of which HOQ is to receive the arriving message, and results in the matrix 66 creating the connection and
 15 passing on the message to the required destination without storing it in the (empty) HOQ.

Finally, if a message with a good CRC arrives at an empty EQ and the node (i.e. the node associated with the egress router which has received the message) is able to accept it, the message is immediately forwarded thus
 20 avoiding another unnecessary buffer read and write.

The passing of messages from the ingress router 60 to the matrix 66 is controlled by a credit protocol. This allows the ingress router 60 to know at all times whether the matrix 66 is able to accept a new message, without the overhead of a request/reply system.

25 The ingress router 60 maintains a credit counter for each VOQ/HOQ pair, which starts out initialised to the capacity in messages of each empty HOQ. When a message is available for transmission to the HOQ, the state of this

counter is examined. If there is available credit (the counter is non-zero) the message is passed via the serial interface and the credit is decremented by one. When a connection_grant is later received indicating a message has been removed from the HOQ, the credit counter is incremented. If there is
 5 insufficient credit for a new message to be sent to the matrix, the message is stored in the VOQ and sent later when credit becomes available.

This preceding discussion represents only one possible way in which the HOQs can be kept up to date with the VOQs. For example, it be possible in principle for each HOQ to request another cell from the respective VOQ
 10 whenever the HOQ transmits a cell through the matrix. However, this variation is not presently preferred since the possibility of signal back from the switch input to the ingress router is not part of the architecture previously described.

A simple extension to this scheme within the scope of the present invention would allow the system to cope with different sized messages. Instead of the
 15 credit counter simply counting whole messages, it could count message words. The counter would then be decremented or incremented by the number of words for the message being added or removed, and the criterion for being able to add a new message would be that the counter would not go negative following the decrement.

20 The available credit at any time is the ingress router's 60 view of the uncommitted spare capacity in that HOQ in the matrix 66.

In the system described above with relation to Figures 1 and 2 only, messages are deleted from their VOQ when they are forwarded to the matrix. In the present embodiment, however, messages are retained in the VOQ until
 25 the egress router 64 reports successful receipt in order to support the reliable delivery protocol. The occupied area of the VOQ can thus conceptually divided into three areas (any or all of which can of course be empty):

- W – This area contains messages waiting for transmission to the HOQ. This area is only occupied when there is currently no credit for further messages to the corresponding HOQ.

- H – This area represents messages that have been forwarded to
5 the HOQ and are waiting there for onward transmission to their destination.

- A – This area contains messages that have been forwarded from the HOQ but for which no response (ack or nack) has yet been received.

Note that although H and A are shown as separate areas of the VOQ, these areas are conceptual: as far as the ingress router 60 is concerned, the areas
10 H and A constitute a single area of cells which have been transmitted to the matrix 66 already. It would be possible for the router to track this boundary through the connection_grant signals, but in practice this is unnecessary and is not done.

Each arriving message is checked for a correct CRC in the egress router 64,
15 and an ack (good) or nack (bad) response generated. A sequence number within the message header is also checked to guard against lost messages – if the sequence number of the latest message is not contiguous with the previous message, a nack response for any missing message is generated before the response for the current message. Responses generated by the
20 egress router 64 are sent to the master 68, which then routes them back to the originating ingress router 60.

In the ingress router 60, an arriving response should always refer to the message at the head of the appropriate VOQ (this is checked by tagging messages and responses with sequence numbers). The message is removed
25 from the VOQ: if the response is ack, the message has been correctly transferred and is therefore abandoned. If the response is nack the message is requeued into the tail of the VOQ and treated as if it were a new message.

The depth of each HOQ buffer is selected to allow operation at the full bandwidth over the credit protocol filling from the VOQs, and to provide sufficient messages to allow the scheduling and arbitration to proceed efficiently under high load conditions.

- 5 For the former requirement, consider a system where the message buffers are empty and a full bandwidth stream of messages starts arriving from the Node. For this message stream to keep flowing out to the HOQs without being queued in the VOQ region W, a new message must never encounter a zero credit. Consider the delay between the first message arrival from the Node
10 and the return of credit for that message:

	connection_request generation:	5 ns
	Router-Master serial control interface	20 ns
	Master arbitration	10 ns
	connection_grant generation	5 ns
15	Master-Router serial control interface	20 ns
	Decode and credit restoration	5 ns
	Total	65 ns

- If at full bandwidth messages arrive in the Router to be sent on this channel every 10ns, this indicates that the HOQ should hold a minimum of 7
20 messages to avoid lack of credit throttling the message flow. In practice of course, the master arbitration could take considerably longer than 10ns due to port contention. Extra HOQ space would defer the onset of flow throttling in such a situation.

To provide sufficient messages to allow the scheduling and arbitration to proceed efficiently under high load conditions, consider a system where the HOQ has filled and the master 68 starts arbitrating a continuous series of connections. Assuming there is a backlog of messages waiting in the VOQ region W, the HOQ should contain enough messages to satisfy the connections until returning credit restarts the flow from the VOQ. Consider the delay between connection generation in the master and new connection_requests arriving from the router:

	connection_grant generation	5 ns
10	Master-Router serial control interface	20 ns
	Decode and credit restoration	5 ns
	Message extraction from VOQ region W	10 ns
	connection_request generation:	5 ns
	Router-Master serial control interface	20 ns
15	Total	65 ns

If connections are generated every 10 ns, this implies that the HOQ should contain a minimum of 7 messages to avoid any interruption to the connections while waiting for new messages from the Router.

Figures 4 to 8 show the sequence and timing of operations in the components of the embodiment. The latency through the embodiment has been defined from message valid in ingress router 60 to message valid and checked in egress router 64. The latency through the serial links is detailed in Figure 4. We will now describe in detail the time taken to perform various operations.

1. Fast Message Transfer

If the master 68 receives a connection request and detects that there is no contention for the ingress and egress ports involved in that request (HOQ and arbiter status), then the master 68 can bypass the arbitration phase (which is
5 there to resolve contention fairly) and immediately grant the connection. This “fast message transfer” feature reduces the message transfer latency when the embodiment is under condition of low load or when it is supporting stochastic, non-contentious flows.

In a fast message transfer, the HOQ routing data is not sent to the matrix
10 device over the Master-Matrix interface since the message does not need to be stored in a HOQ.

The timescales for the operation are as set out in Fig. 5. In the absence of contention for destination ports, the embodiment supports the full port bandwidth across all ports with the 55ns “fast message transfer” latency as
15 shown above. The embodiment is a strictly non-blocking fabric, so there is no input or internal contention.

2. Arbitrated Message Transfer

Fig. 6 shows the timing of a “normal” message transfer where there is some contention in the fabric, but where the requested message is forwarded with
20 no extra queuing delay. In the event of contention for an output port (two or more messages in the fabric destined for delivery to the same output node), the limited bandwidth of the router-node interface forces all but one of the messages to be queued in the HOQ buffers. This queuing due to collisions
25 between messages appears as an increase in the average latency through the fabric. The magnitude of this increase depends on the traffic patterns of the application (probability of message collisions).

Fig. 7 shows the average message latency through the embodiment a 16 port TSI assuming that all ports are sending to all other ports with equal probability and with random inter-message gaps. The chart illustrates that the embodiment's performance is close to the optimal behaviour of an M/M/1 queue (that is, a Mark of/Mark of queue, a term of art which refers to a single server Q with Poisson distributed arrival rates and Poisson distributed service rates) queue, particularly compared to a simple FIFO queued fabric (no VOQs).

Note that this chart does not illustrate the effect of the fast message transfer described above (e.g. in relation to Fig. 5) which will further reduce the average latency at low loads.

It should also be noted that the increase in message latency under conditions of high loading is not a feature of the fabric, but is caused by output contention, i.e., to alleviate this effect, a node would have to be capable of accepting (and processing) messages at a faster rate (~20% faster) than an individual node could issue messages. In practice this would only serve to move the point of contention further down the dataflow without necessarily improving the overall system performance.

Router ack/nack latency defines the minimum depth of the VOQs required in order to maintain a full bandwidth flow between two nodes. Fig. 8 shows the normal ack/nack latency is 115 ns. With 10ns messages, this indicates an absolute minimum VOQ depth of 12 messages for the H & A regions. The size of the W region is determined by the latency of resuming a paused interface from the Node.

Although the invention has been described above in relation to a single embodiment only, many variations are possible within the scope of the invention. For example, the present invention is not limited to multi-channel transmission. Furthermore, the present invention is not limited to data

transmission between data processors, but rather may be used in any digital communication system.

Also, although the invention has been described above in relation to cells which are each sent only to a single node, the present invention is applicable
5 also in the case of multicast signals. For example, a cell which is to be transmitted to more than one egress router may be divided by the ingress router into a plurality of cells each for transmission to a single egress router. Similarly, cells which are to be sent to multiple outputs associated with a single egress router may contain this information in their headers, so that the
10 egress router may transmit them accordingly.

Similarly, although the cells of the present invention are usually of equal length, some of the field of a given cell may be "invalid", in the sense that they are not used to carry useful information.

Claims

1. A data switching device having a plurality of ingress routers, a plurality of egress routers, a switching matrix and a connection controller,
the switching matrix having input ports connected to respective said
5 ingress routers and output ports connected to respective said egress routers, and controlled by the controller to form connections between pairs of input and output ports;
each ingress router including one or more virtual output queues for each egress router, each virtual output queue being arranged to store fixed
10 length cells having a header defining the egress router to be used in the switching matrix connection;
each input port of the switching matrix including for each virtual output queue in the ingress router connected to that input port a respective head of queue buffer, each buffer being arranged to receive a replication of at least
15 one cell in the corresponding virtual output queue;
wherein, upon the switching matrix forming a connection between a given input port and output port, the switching matrix transmits from that input port to that output port a cell from one of the one or more corresponding head of queue buffers; and
20 upon error free receipt by an egress router of a cell from one of the virtual output queues of one of the ingress routers, a receipt signal is transmitted to that ingress router, that ingress router storing the cell until receiving the receipt signal.
2. A device according to claim 1 in which, upon receipt of a new cell by
25 one of the ingress routers, the cell is stored in an appropriate said virtual output queue of the ingress router, and, if a credit count, indicative of the number of free cells of the corresponding head of queue buffer, is not zero, a

replication of the cell is transmitted to that head of queue buffer and a connection request is transmitted to the controller.

3. A device according to claim 2 in which, upon the controller causing a connection to switch a cell of an ingress router through the switching matrix, a connection grant signal is transmitted to that ingress router, and increments the credit count by one.

4. A device according to claim 2 or claim 3 in which, upon said replication of the cell to the head of queue buffer, the respective credit count is decremented by one.

5. A device according to any preceding claim in which the virtual output queues are segregated into two areas, a first area containing cells waiting for replication to the corresponding head of queue buffer, and a second area containing cells replicated to the head of queue buffer, and, upon determining that there is at least one cell in the first area and that the number of free cells of the corresponding head of queue buffer is not zero, a replication of at least one cell in the first area is transmitted to that head of queue buffer, the cell is transferred to the second area, and a connection request is transmitted to the controller.

6. A device according to any preceding claim in which each cell is associated with a priority level, said virtual output queues comprising a virtual output queue for cells of each respective priority level, said controller determining, in tandem with which pairs of input and output ports will be connected, the priority level of the cell to be transmitted between them.

7. A device according to any preceding claim in which the controller determines whether any given one of the cells in the virtual output queues can be transmitted between the appropriate pair of input and output ports without preventing the transmission of a cell in a virtual output queue between

another pair of input and output ports, and in this case causes that given cell to be transmitted.

8. A device according to any preceding claim in which each egress router is arranged to detect that a cell transmitted by the switching matrix has not
5 been received correctly, and in this case transmits a re-transmission request to the corresponding ingress router.

9. A device according to claim 5 and claim 8 in which, upon receiving the re-transmission request, the ingress router transfers the corresponding cell in
10 the second area into the first area, and transmits a corresponding connection request to the controller.

10. A method of operating a data switching device having a plurality of ingress routers, a plurality of egress routers, a switching matrix and a connection controller, and

15 the switching matrix having input ports connected to respective said ingress routers and output ports connected to respective said egress routers, and controlled by the controller to form connections between pairs of input and output ports;

the method comprising the steps of:

20 maintaining at each ingress router one or more virtual output queues for each egress router, each virtual output queue being arranged to store fixed length cells having a header defining the egress router to be used in the switching matrix connection;

maintaining at each input port of the switching matrix for each virtual
25 output queue in the ingress router connected to that input port a respective head of queue buffer, each buffer being arranged to receive a replication of at least one cell in the corresponding virtual output queue,

upon the switching matrix forming a connection between a given input port and output port, transmitting from that input port to that output port a cell from one of the one or more corresponding head of queue buffers, and

5 upon error free receipt by an egress router of a cell from one of the virtual output queues of one of the ingress routers, transmitting a receipt signal to that ingress router, that ingress router storing the cell until receiving the receipt signal.

11. A method according to claim 10 in which, upon receipt of a new cell by one of the ingress routers, the cell is stored in an appropriate said virtual output queue of the ingress router, and, if a credit count, indicative of the number of free cells of the corresponding head of queue buffer, is not zero, a replication of the cell is transmitted to that head of queue buffer and a connection request is transmitted to the controller.

12. A method according to claim 11 in which, upon the controller causing a connection to switch a cell of an ingress router through the switching matrix, a connection grant signal is transmitted to that ingress router, and increments the credit count by one.

13. A method according to claim 11 or claim 12 in which, upon said replication of the cell to the head of queue buffer, the respective credit count is decremented by one.

14. A method according to any of claims 10 to 13 in which the virtual output queues are segregated into two areas, a first area containing cells waiting for replication to the corresponding head of queue buffer, and a second area containing cells replicated to the head of queue buffer, and, upon determining that there is at least one cell in the first area and that the number of free cells of the corresponding head of queue buffer is not zero, a replication of at least one cell in the first area is transmitted to that head of queue buffer, the cell is

transferred to the second area, and a connection request is transmitted to the controller.

15. A method according to any of claims 10 to 14 in which each cell is associated with a priority level, said virtual output queues comprising a virtual
5 output queue for cells of each respective priority level, said controller determining, in tandem with which pairs of input and output ports will be connected, the priority level of the cell to be transmitted between them.

16. A method according to any of claims 10 to 15 in which the controller determines whether any given one of the cells in the virtual output queues can
10 be transmitted between the appropriate pair of input and output ports without preventing the transmission of a cell in a virtual output queue between another pair of input and output ports, and in this case causes that given cell to be transmitted.

17. A method according to any of claims 10 to 16 in which each egress
15 detects that a cell transmitted by the switching matrix has not been received correctly, and in this case transmits a re-transmission request to the corresponding ingress router.

18. A method according to claim 14 and claim 17 in which, upon receiving
20 the re-transmission request, the ingress router transfers the corresponding cell in the second area into the first area, and transmits a corresponding connection request to the controller.

19. A switching device substantially as described herein with reference to Figures 2 to 6 and 8.

25 20. A method of operating a data switching device substantially as described herein with reference to Figures 2 to 6 and 8.

THIS PAGE BLANK (USPTO)




Fig. 1

THIS PAGE BLANK (USPTO)

217




Fig. 2

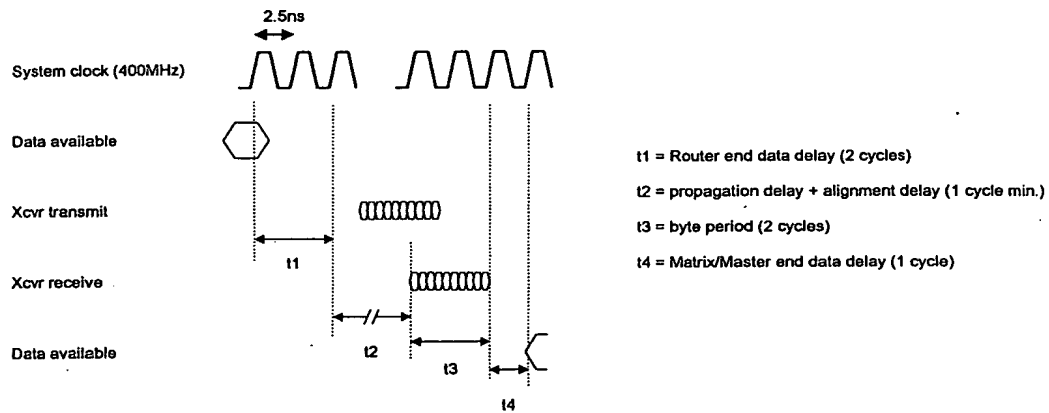



Fig. 4

THIS PAGE BLANK (USPTO)



THIS PAGE BLANK (USPTO)

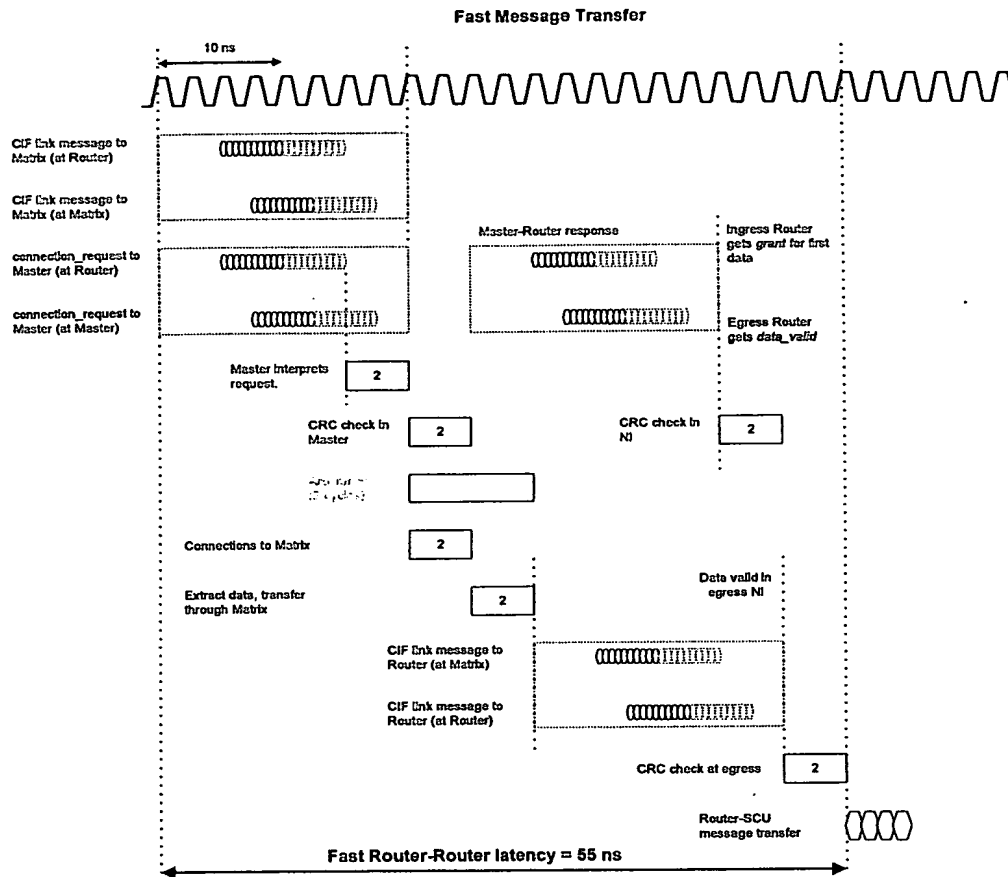


Fig. 5

THIS PAGE BLANK (USPTO)

517

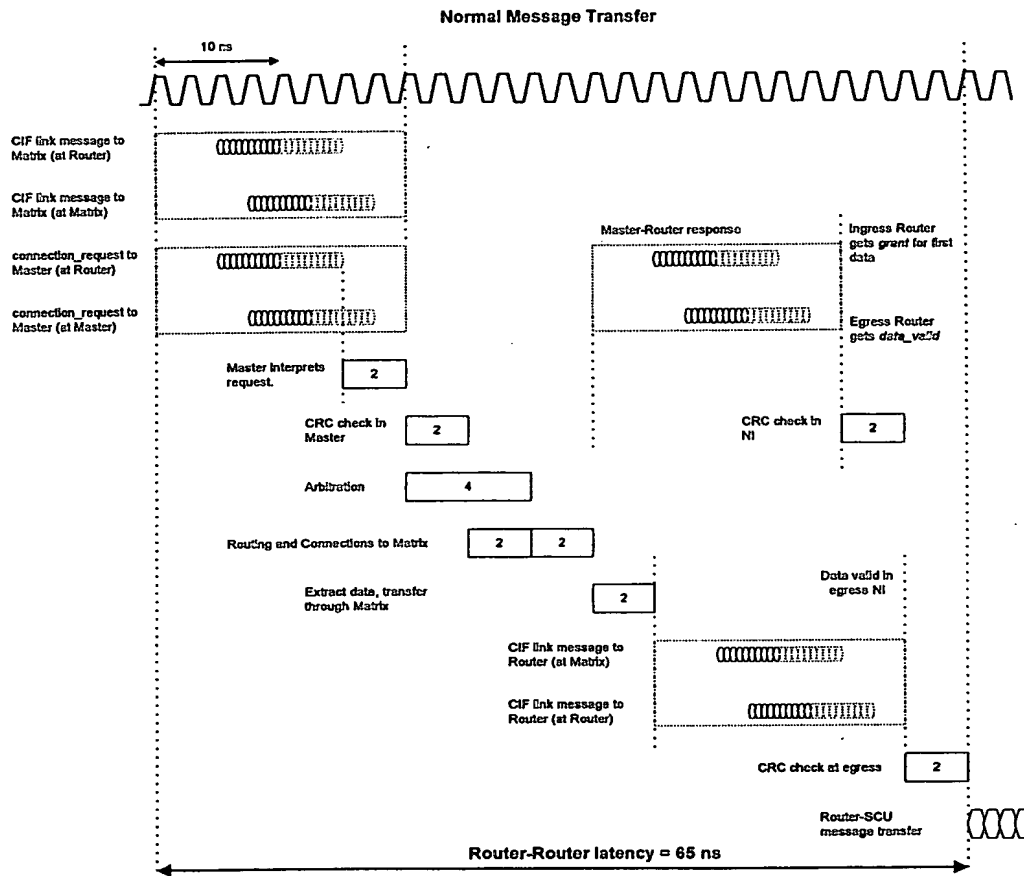


Fig. 6

THIS PAGE BLANK (USPTO)

6/7

Loaded Latency




Fig. 7

THIS PAGE BLANK (USPTO)

7/7

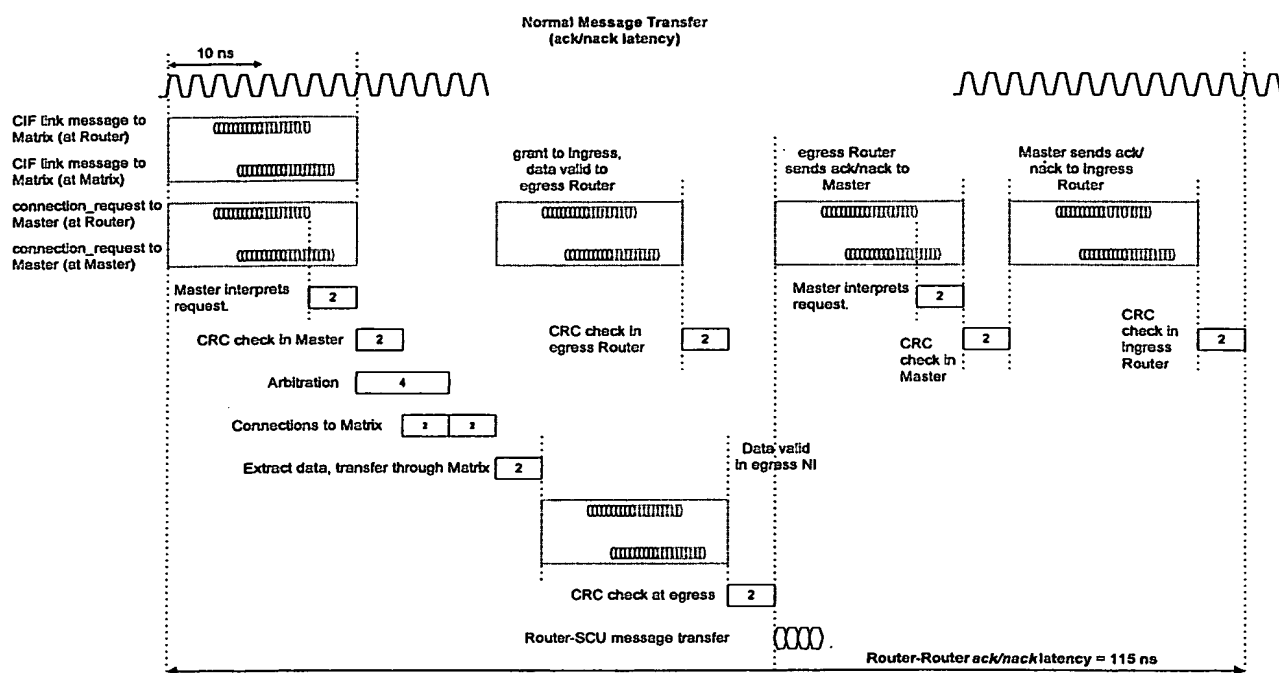


Fig. 8

THIS PAGE BLANK (USPTO)